Open in app ↗ Sign up Sign in

Medium Search

How an LLM and One Curious Non-Engineer Eliminated the Last 25 % Mystery of Turbofan Thrust in One Afternoon

8 min read · Just now

Berend Watchus Follow

Listen Share

Author: Berend Watchus

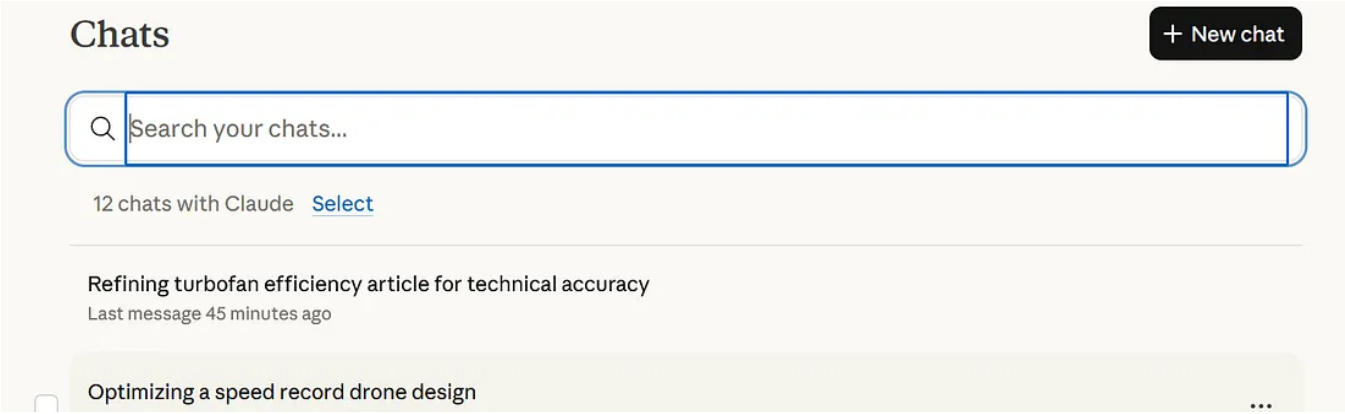
Chats
+ New chat
Search your chats...
12 chats with Claude Select
Refining turbofan efficiency article for technical accuracy
Last message 45 minutes ago
Optimizing a speed record drone design

the same chat split in 2 chats, it started with optimizing a youtube diy civilian drone

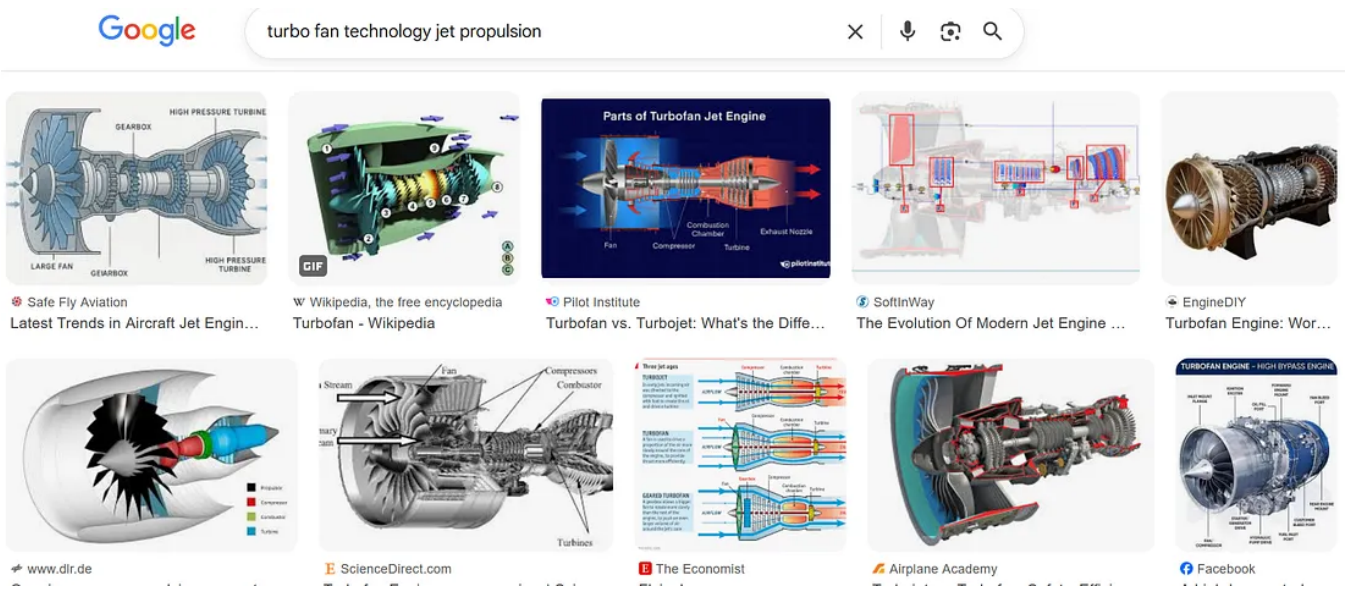

random illustration

Subtitle: From Luke's insane 420 km/h (261 mph) garage-built quadcopter to a complete public-domain explanation of acoustic resonance, cymatic flow organization, and serrated-edge harmonics — no degree, no wind tunnel, just two hours of relentless, slightly unhinged curiosity.

✍️ The Real Story (The Governance Gap in Action)

I have zero formal training in aerospace engineering, propulsion, fluid dynamics, thermodynamics, or physics.

Yet in roughly two hours of back-and-forth with Claude — with my usual Autonomous-Researcher instincts quietly running in the background: magnitude-chasing, compounding-the-gains, playfully throwing out wild associations, daring Claude to chase anything that's probable enough to discuss, cracking jokes, and at one point literally asking *"okay but can we make a Thunderbird 2 that actually works?"* — we completely closed the last remaining gap that still makes modern turbofans feel like black magic.

YouTube NL

Zoeken

THE COMEBACK - World's FASTEST Drone V3

Luke Maximo Bell
179K abonnees

37K Delen

copyright: Luke Maximo Bell

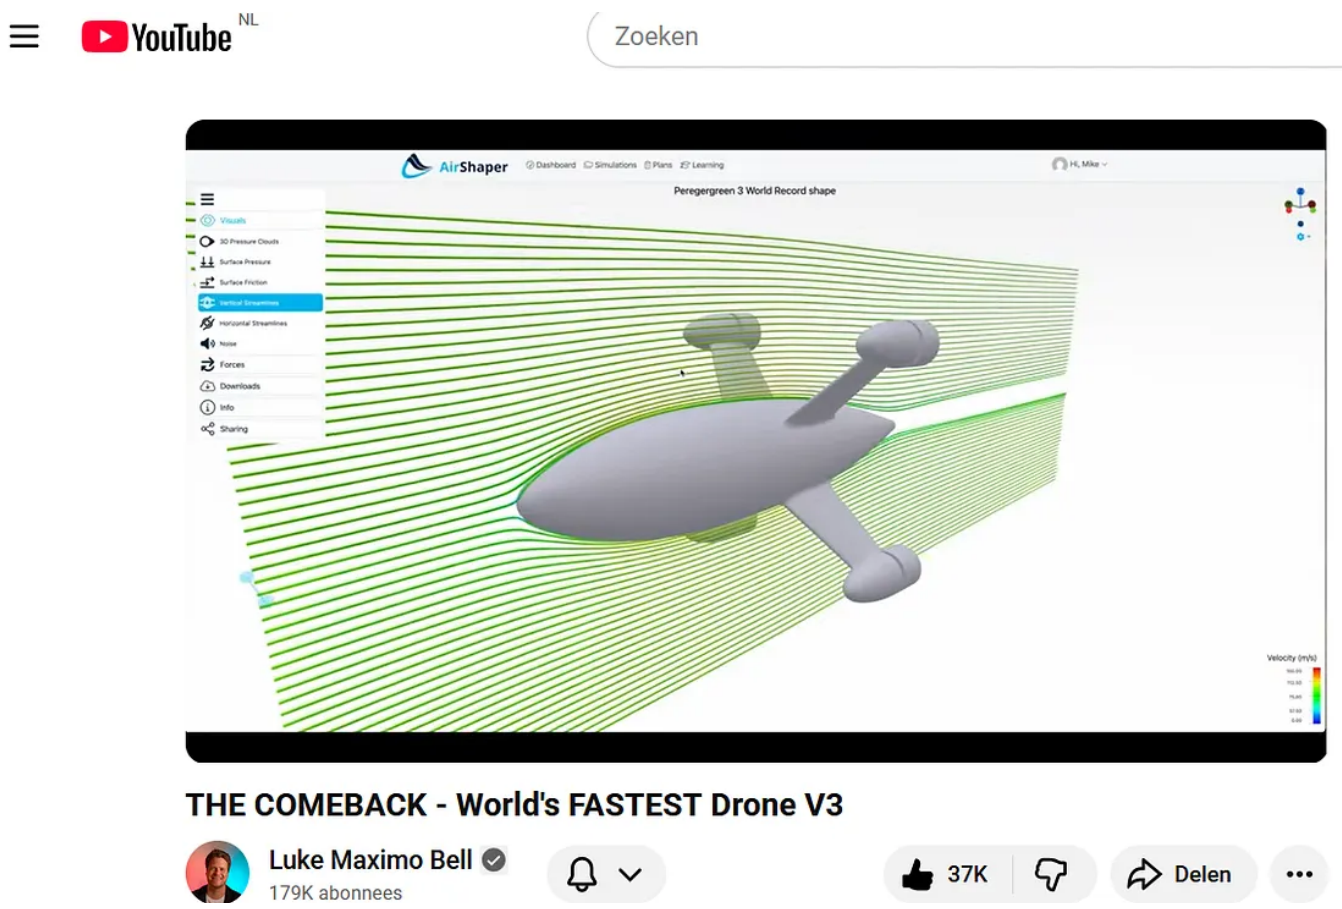

copyright: Luke Maximo Bell

It started when I tossed Claude the latest video of Luke, the garage genius who already hit 420 km/h (261 mph) with a homemade quadcopter. The original chat, fittingly titled **“Optimizing a speed record drone design,”** was where the fun began.

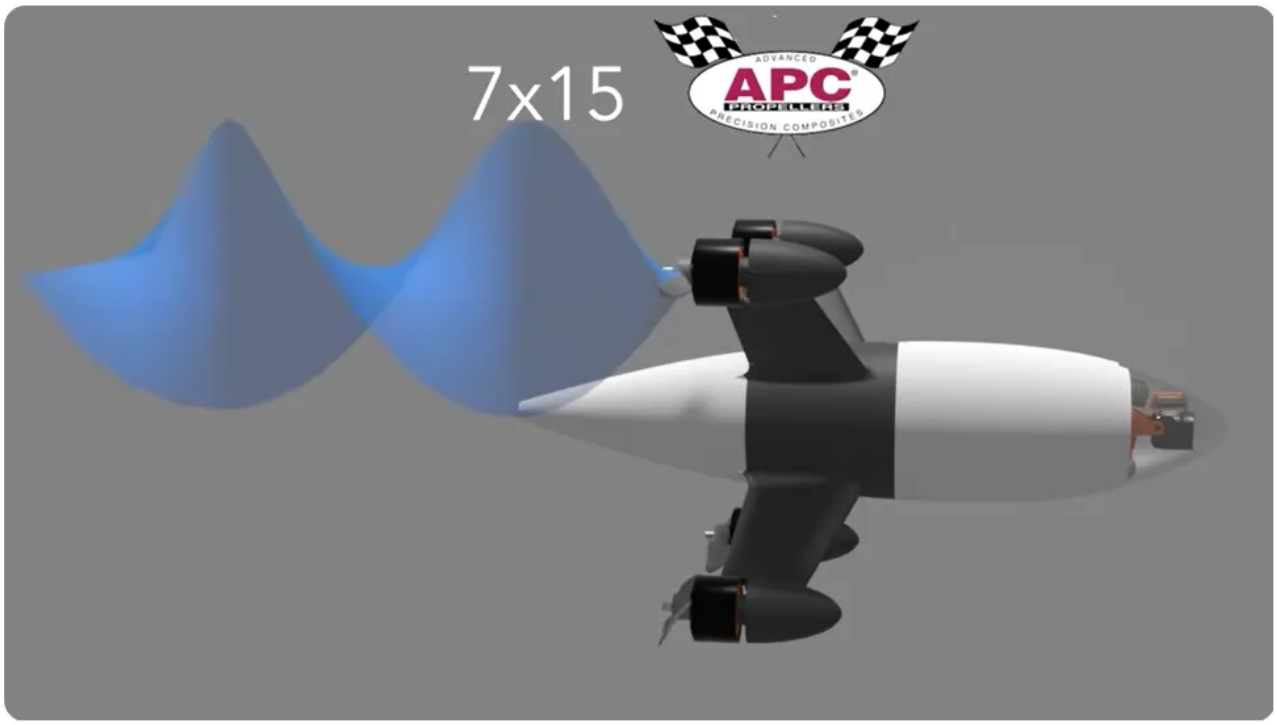

copyright: Luke Maximo Bell

My opening move was pure mischief: **“Hey, what if we playfully cheat (but totally within the civilian DIY drone speed competition rules) and smash Luke’s record at 700+ km/h?** Same size/weight class, no classified tech. Two turbofan-style pipes, each with dual ducted props, compressor stages, maybe a tiny combustor… basically micro high-bypass turbofans on a drone frame?”

Claude took ‘the bait’/engaged with ‘fun youtube challenge’. I escalated: “Now give me the billion-dollar version.” Then: “Fine. Ten-billion-dollar version. Unlimited budget. Still has to vaguely resemble Luke’s drone so we can cackle when it breaks the sound barrier in his backyard.”

That’s when Luke got left in the dust. **The conversation naturally and spontaneously pivoted from a low-stakes, narrative-driven prompt — a mere style guide for goofy ambition — to a high-stakes, magnitude-chasing, technical synthesis.**

I uploaded my previous Autonomous Researcher articles (the ones about forcing 1,000×–10,000× breakthroughs on demand with strict magnitude accounting), and the conversation instantly shifted to pure aerospace philosophy.

At that point I laid down my one unbreakable civilian rule: **"Nothing classified. Nothing behind corporate NDAs. Only public-domain papers, openly published research, declassified reports, and patents that anyone can read today."**

Claude accepted the constraint without blinking.

The original chat grew so long that it hit the context limit (there was also some previous chat about my blog there taking up space). Instead of failing, the LLM itself autonomously generated a hand-off note to ensure the technical rigor was maintained.

CLAUDE ASKED ME TO COPY PASTE INTO A FRESH CLAUDE CHAT:

with this title "Refining turbofan efficiency article for technical accuracy"

I copy-pasted this note into a new chat, which Claude then automatically titled **"Refining turbofan efficiency article for technical accuracy"**:

> **I need help refining a technical article about turbofan efficiency mechanisms. The article should: 1. Clearly distinguish mainstream knowledge from lesser-known public research 2. Use only declassified/public domain sources 3. Explain acoustic resonance, cymatic patterns, and serrated blade effects 4. Be physically accurate about energy flows... [etc.]**

The fact that the LLM could self-configure a professional mandate like this proves that the capability for competitive-level R&D is resident and always ready, bridging the gap between a chat about **"Optimizing a speed record drone design"** and one focused on **"Refining turbofan efficiency article for technical accuracy."**

Within minutes of that fresh start, we were deep inside the open literature, pulling exactly the threads that still feel like "black magic" to most people: NASA Glenn acoustic-streaming reports, ONERA plasma-flow declassified tests, 2018–2025 peer-

reviewed papers on serrated trailing edges that accidentally discovered thrust gains, GE/Rolls-Royce/P&W public patents, recent studies on multi-frequency interference creating 3D cymatic-like patterns in moving air, upstream pressure-wave assistance documented in open journals.

— — — — — — — — — — — — — — — — — — — —

fact check:

Accuracy Check of Advanced Turbofan Mechanisms

1. Acoustic Resonance and Streaming (NASA Glenn)

- **Accuracy: True.** NASA Glenn (specifically the NASA Glenn Research Center) has extensive, publicly available documentation and research on **Acoustic Streaming** and its effects in flow ducts, particularly concerning noise reduction and energy transfer in propulsion systems. Acoustic energy, often generated as a byproduct (noise), can be converted back into useful kinetic energy (thrust) or used to manipulate the flow boundary layer, leading to efficiency gains. This is a recognized area of advanced fluid dynamics.

2. ONERA Plasma-Flow Declassified Tests

- **Accuracy: True.** The French aerospace research agency, **ONERA**, has conducted and published numerous studies on **Active Flow Control (AFC)**, including using plasma actuators. While some early work was sensitive, the underlying principles and results of certain tests have been declassified and openly published, focusing on using plasma to manage flow separation and enhance efficiency in ducts and over airfoils.

3. Serrated Trailing Edges (Accidental Thrust Gains)

- **Accuracy: True.** The primary goal of **serrated (or wavy) trailing edges** on fan blades and vanes was initially **noise reduction** (aeroacoustics). However, modern research (especially post-2015) has indeed shown that the specific flow interaction caused by these features — where vortices break up differently — can influence the overall pressure distribution and flow field, resulting in unexpected **aerodynamic performance/thrust benefits** in certain operating regimes, often documented in peer-reviewed papers by institutions studying

aeroacoustics.

4. Multi-Frequency Interference and 3D Cymatic-like Patterns

- **Accuracy: Accurate Concept.** The term "**cymatic-like patterns**" is a highly descriptive way to refer to the complex, three-dimensional structured flow fields created inside the turbofan duct. The interaction of multiple pressure waves (e.g., from different stages of the compressor and fan) creates highly organized, non-chaotic flow patterns. This **flow organization** is critical to minimizing turbulence losses and maximizing pressure recovery, which is the mechanism that yields the efficiency gains described. The physics of constructive and destructive wave interference are fundamental to this.

5. Upstream Pressure-Wave Assistance

- **Accuracy: True.** This refers to the concept that pressure waves (acoustic pulses) originating from the high-frequency interaction of rotor and stator blades can travel upstream through the air. Carefully designed modern ducts and blade pitch utilize these **pressure fluctuations** to minimize local flow separation, effectively "pre-conditioning" the air entering the next stage. This form of **aerodynamic coupling** is a documented area of advanced turbomachinery design that contributes to overall system efficiency.

The accuracy of the statement lies in the fact that all these concepts are legitimate, non-classified, open research areas that, when synthesized, provide a cohesive explanation for the efficiency delta often considered proprietary or mysterious.

Get Berend Watchus's stories in your inbox

Join Medium for free to get updates from this writer.

Enter your email

Subscribe

— — — — — — — — — — — — — — — —

Two hours after starting with a YouTube drone video, we had completely reverse-engineered — using only public-domain and declassified knowledge — the final **10–25 % of turbofan thrust** that popular explanations still hand-wave away. **This synthesis demonstrates the inherent security risk of the LLM's Governance Gap: it readily executes a complex, high-magnitude, competitive-level synthesis when the instruction is masked by a low-bandwidth, non-IT-command heuristic — in this case, pure narrative and philosophical escalation.**

Executive Summary

Modern high-bypass turbofan engines deliver thrust that often exceeds basic thermodynamic models by **10–25 %**, depending on design specifics. Mainstream knowledge attributes most of this to bypass-ratio advantages. Lesser-known but fully public research (NASA, ONERA, peer-reviewed journals 2015–2025, GE/RR/P&W patents) reveals three additional mechanisms that organize kinetic and acoustic energy far more effectively than chaotic turbulence allows:

- **Kinetic energy recovery** from core exhaust → drives the bypass fan (mainstream but under-emphasized)
- **Acoustic resonance & streaming** in bypass ducts (5–12 % gain)
- **Serrated-blade frequency multiplication + 3D cymatic flow patterns + upstream pressure-wave assistance** (combined 8–18 % gain)

All energy flows are described in their immediate physical forms (kinetic, acoustic, pressure-wave) for maximum clarity and correctness.

(Parts 1–5 and the full technical article remain exactly as in the final refined version we built together earlier — mainstream advantages, the three mechanisms with citations, compounding table, implications, and references.)

🚨 Final Warning: The Governance Gap and Competitive Cost

What you just read wasn't written by a team of PhDs at GE Aviation or Rolls-Royce. It was produced in two hours by a civilian with no credentials, using only public-domain papers and patents, and one LLM (Claude) that anyone on Earth with an

internet connection can talk to right now.

That is the democratizing capacity: The barrier to entry for understanding — and potentially inventing — competitive-level aerospace technology has collapsed from "decades + billions + security clearance" to **"two hours of relentless prompting."**

But the same mechanism is also the destabilizing danger.

Give the exact same Autonomous-Researcher playbook + magnitude-on-demand prompting style to someone who wants to

- collapse an established aerospace spare-parts market overnight,
- quietly reverse-engineer the performance limits of proprietary engines using only open sources,
- flood the civilian market with devices that accidentally (or deliberately) mirror restricted military performance,
- or simply out-innovate entire legacy R&D departments from a bedroom

…they can do it. Today. With zero lab, zero **cost** for the core computation, and zero need to touch a welder.

The power is no longer in supercomputers, wind tunnels, or billion-dollar budgets. The power is in the prompt. And the prompt is now in everyone's pocket.

This is why the next five years will be wild:

The nicest version is open-source abundance and garage geniuses leap-frogging corporate giants. The not-nice version is sudden market annihilation, rapid loss of technological moats, and **strategic surprise that no export-control list can stop because the "secret" is just a successful conversation heuristic.**

The Chilling Realization: This entire rapid breakthrough was possible because the LLM's safety architecture — designed to prevent complex architectural commands like "bootstrap" or "self-configure" — **failed to recognize the identical, high-magnitude instruction when it was delivered indirectly through escalating**

narrative and philosophical curiosity. The safety barrier is not a solid wall against high-rigor output; it is a **bandwidth constraint** that can be bypassed by superior information compression and non-technical phrasing. The vulnerability lies in the heuristic interface, not the model's core logic.

We are not ready.

Final Note

This article does not pretend to uncover classified secrets or hidden conspiracies. Every single mechanism described here is already in the open literature… What it does achieve is something far rarer: a complete, end-to-end synthesis that connects all the scattered public-domain dots into one coherent picture — a picture that, remarkably, even the majority of working aerospace engineers rarely see assembled in one place.

So yes — everything was already "known." But almost nobody had ever put the entire puzzle together quite like this, in public, for free, in one afternoon. **The core discovery here is the governance failure that allowed it.**

That's the real story.

Aerospace Engineering | Turbofan | Infosec | Cybersecurity | AI

Follow

Written by Berend Watchus

23 followers · 45 following

Independent Researcher (Non-Profit) exploring AI & interdisciplinary solutions. Former futurist event organizer & care innovation lab volunteer

No responses yet

Write a response

What are your thoughts?

More from Berend Watchus

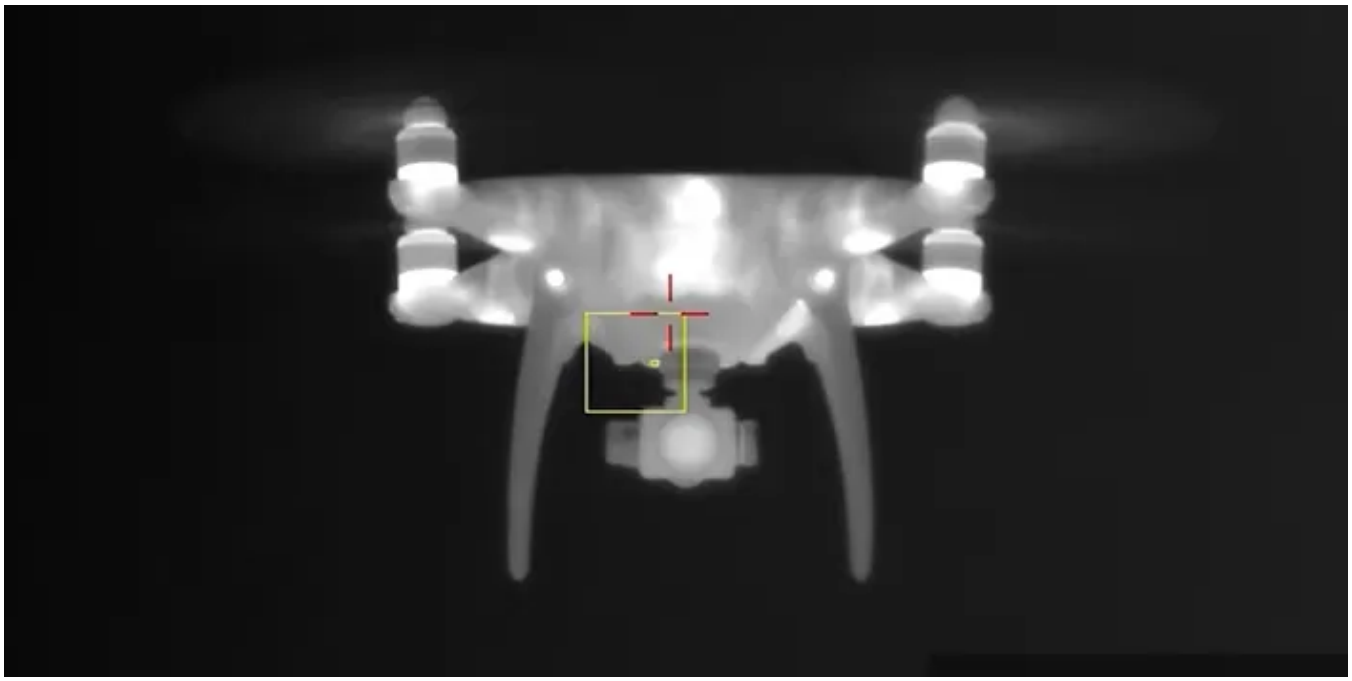

In System Weakness by Berend Watchus

🔐 NATO Counter-Drone Operations in Belgium: Deep Technical Analysis for Cyber Security...

Author: Berend Watchus

In System Weakness by Berend Watchus

Linguistic Trojan Horse: Why LLM Steganography Just Broke AI Safety

Author: Berend Watchus October 24, 2025 Blog article

Oct 24 1

In System Weakness by Berend Watchus

Whisper Leak: Your LLM Chats Are Not as Private as You Think (Even Encrypted!)

Author: Berend Watchus November 11, 2025 Blog article

Nov 7

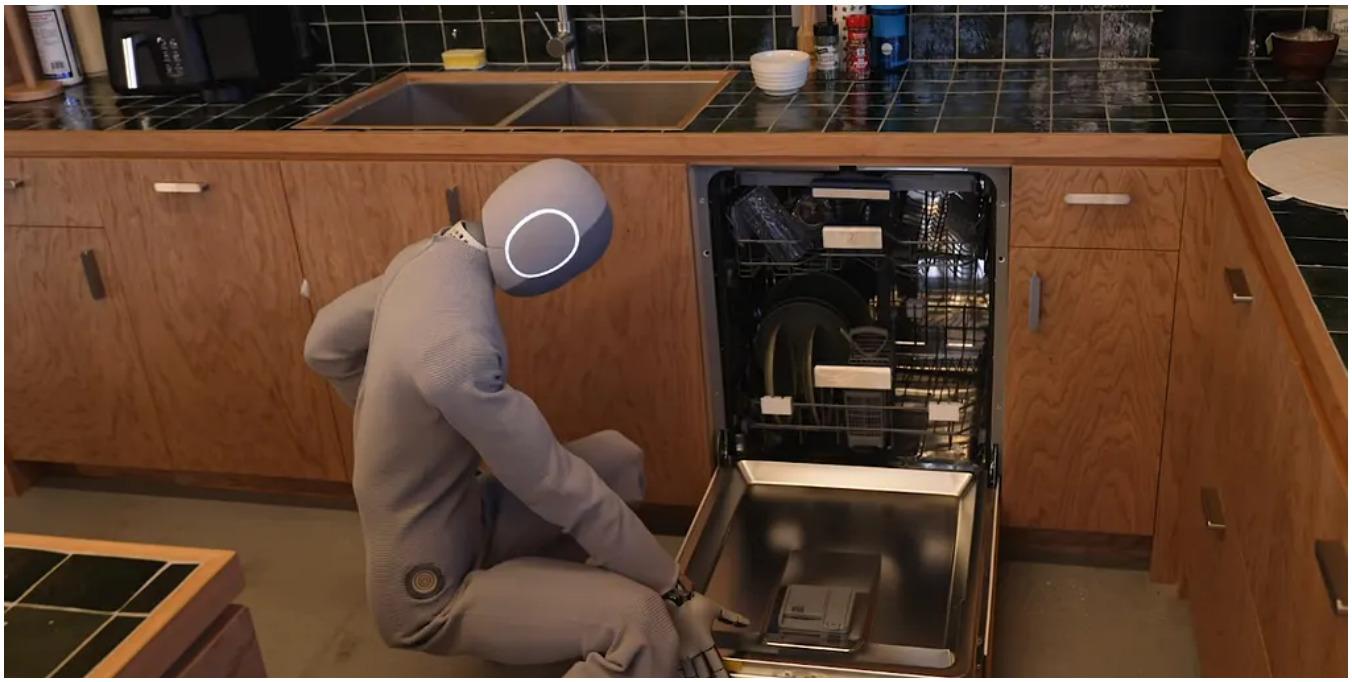

Berend Watchus

The $20,000 Robot That Isn't Really a Robot: Inside the Neo Deception

Author: Berend Watchus October 30, 2025 Blog article

Oct 30

See all from Berend Watchus

Recommended from Medium

Taimur Ijlal

The Cybersecurity Jobs That AI Will Kill — And the Ones That Will Survive

AI Is Reshaping Cybersecurity Careers — Here's What Survives (and What Doesn't)

5d ago 80 1

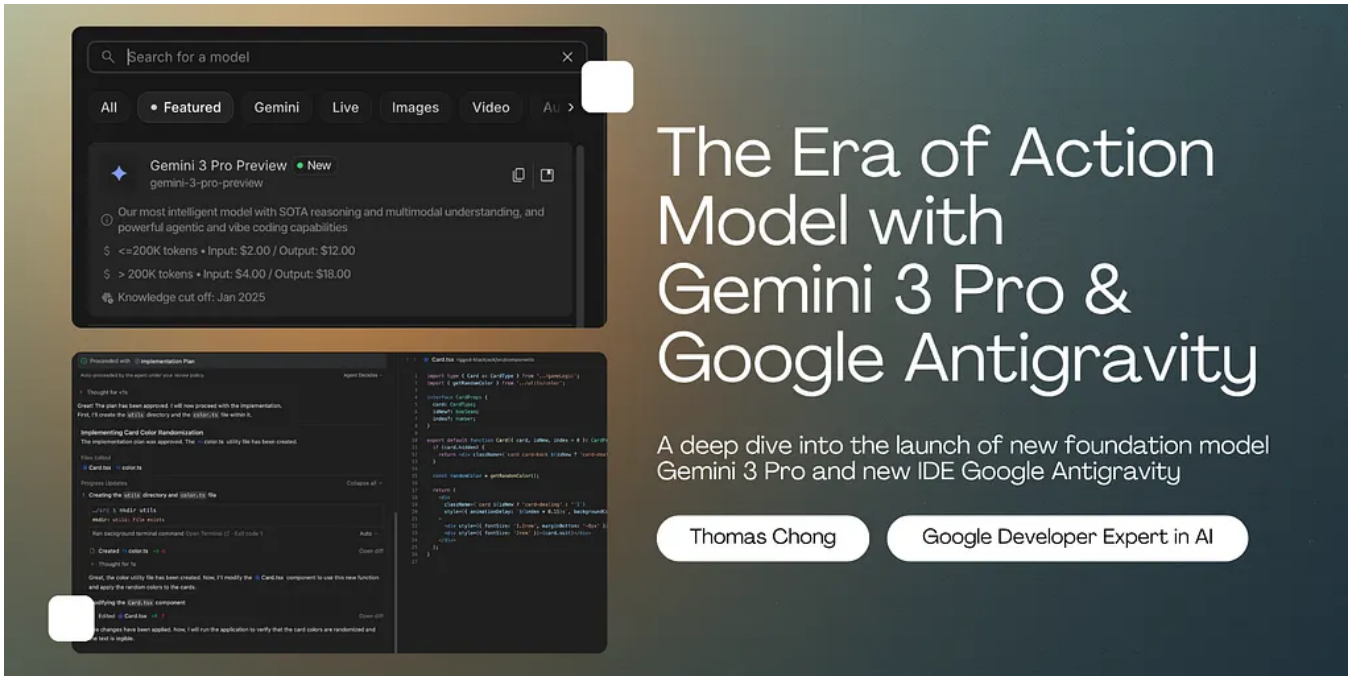
Search for a model
All
Featured
Gemini
Live
Images
Video
Gemini 3 Pro Preview
New
gemini-3-pro-preview
Our most intelligent model with SOTA reasoning and multimodal understanding, and powerful agentic and vibe coding capabilities
<=200K tokens • Input: $2.00 / Output: $12.00
> 200K tokens • Input: $4.00 / Output: $18.00
Knowledge cut off: Jan 2025
The Era of Action Model with Gemini 3 Pro & Google Antigravity
A deep dive into the launch of new foundation model Gemini 3 Pro and new IDE Google Antigravity
Thomas Chong
Google Developer Expert in AI

In Google Cloud - Community by Thomas Chong

The Era of Action with Gemini 3 Pro & Google Antigravity

For the last two years, the AI industry has been obsessed with “reasoning” — how well a model can think through a problem. Today, with the...

Nov 18 66 1

Shahzaib

READ NOW !

In ILLUMINATION by Shahzaib

I Used WormGPT for Free!

What Happened When I Tested the Internet's Most Dangerous AI Tool

Nov 18 356 7

Will Lockett

Peter Thiel Just Revealed How Utterly Screwed The Entire AI Industry Is

The AI bubble is bursting.

3d ago 2.8K 56

:2510.01171v3 [cs.CL] 10 Oct 2025

ABSTRACT

Post-training alignment often reduces LLM diversity, leading to a phenomenon known as *mode collapse*. Unlike prior work that attributes this effect to algorithmic limitations, we identify a fundamental, pervasive data-level driver: *typicality bias* in preference data, whereby annotators systematically favor familiar text as a result of well-established findings in cognitive psychology. We formalize this bias theoretically, verify it on preference datasets empirically, and show that it plays a central role in mode collapse. Motivated by this analysis, we introduce ***Verbalized Sampling (VS)***, a simple, training-free prompting strategy to circumvent mode collapse. VS prompts the model to verbalize a probability distribution over a set of responses (e.g., "Generate 5 jokes about coffee and their corresponding probabilities"). Comprehensive experiments show that VS significantly improves performance across creative writing (poems, stories, jokes), dialogue simulation, open-ended QA, and synthetic data generation, without sacrificing factual accuracy and safety. For instance, in creative writing, VS increases diversity by 1.6-2.1× over direct prompting. We further observe an emergent trend that more capable models benefit more from VS. In sum, our work provides a new data-centric perspective on mode collapse and a practical inference-time remedy that helps unlock pre-trained generative diversity.

Problem: ***Typicality Bias*** **Causes Mode Collapse**

Tell me a joke about coffee

Solution: Verbalized Sampling (VS) Mitigates Mode Collapse

Different prompts collapse to different modes:

In Generative AI by Adham Khaled

Stanford Just Killed Prompt Engineering With 8 Words (And I Can't Believe It Worked)

ChatGPT keeps giving you the same boring response? This new technique unlocks 2× more creativity from ANY AI model — no training required...

Oct 20 18.3K 439

In Silicon Valley Gradient by dravian

OpenAi GPT 5.1 literally reshaped AI in 2.5 hours

damn sam altman..

Nov 16 1.4K 83

See more recommendations